

What is claimed is

1. A method for classifying a plurality of data flows in a router comprising the steps of:

5 partitioning a ternary content addressable memory (TCAM) into at least a first partition and a second partition; said first partition includes indices having highest priority ranging from a lowest index to a partition index and said second partition includes indices having lowest priority ranging from a highest index to said partition index;

10 loading one or more first flow TCAM entries of a first of said plurality of data flows into said first partition in a predetermined order;

 loading one or more second TCAM entries of a second of said plurality of data flows into said second partition in a predetermined order;

15 setting bit values of a corresponding mask for each of said first TCAM entries and said second TCAM entries such that bits of said respective first TCAM entries and said second TCAM entries are individually masked by said masks; and

20 comparing a prefix comprising predetermined packet header information of an incoming packet to said loaded one or more first TCAM entries and one or more second TCAM entries such that a matching said one or more first TCAM entries subsumes any matching said one or more second TCAM entries.

2. The method of claim 1 where said first plurality of data flows are MPLS or IP-VPN flows.

25 3. The method of claim 1 wherein said second plurality of data flows are policy based routing flows.

30 4. The method of claim 3 wherein said policy based routing flows are access control list (ACL) flows.

5. The method of claim 3 wherein said policy based routing flows are traffic manager flows.

6. The method of claim 1 further comprising the steps of:
5 maintaining a flow index space having entries corresponding to said TCAM; and
determining said predetermined order of said first TCAM entries and said
predetermined order of said second TCAM entries in said flow index space before said
steps of loading said one or more first TCAM entries and loading said one or more
second TCAM entries.

7. The method of claim 2 wherein said MPLS or IP-VPN flows are classified by
connection index (CIX) and destination address (DA), CIX only or DA only.

8. The method of claim 7 wherein said first partition is divided into a first portion
15 and a second portion, said first portion includes indices having highest priority ranging
from said lowest index to a gMaxCixDaFix index and said second portion includes
indices having lowest priority ranging from a gMinDaOnlyFix index to said partition
index and said flows classified by CIX and DA and CIX only are assigned to said first
portion and said flows classified by DA only are assigned to said second portion.

9. The method of claim 8 wherein said flows classified by CIX and DA and said
flows classified by CIX only are loaded into said first portion starting at said
gMaxCixDaFix index and said flows classified by said DA only are loaded in said second
portion starting at said gMinDaOnlyFix index.

10. The method of claim 9 wherein said first partition is divided into a first
portion and a second portion, said first portion includes indices having highest priority
ranging from said lowest index to a gMaxCixDaFix index and said second portion
includes indices having lowest priority ranging from a gMinDaOnlyFix index to said
30 partition index, said gMaxCixDaFix index being adjacent to said gMinDaOnlyFix index

and a free entry being positioned below said gMinDaOnlyFix index in said second portion and further comprising the steps of:

(a) moving a flow classified by DA only at said gMinDaOnlyFix index to said free entry;

(b) assigning a flow classified by CIX and DA or CIX only to said gMinDaOnlyFix index after step (a);

(c) setting the gMaxCixDaFix index at the index assigned to said flow classified by CIX and DA or CIX only in step (b); and

(d) setting the gMinDaOnlyFix index immediately after the index assigned to said flow classified by CIX and DA or CIX only in step (b).

11. The method of claim 10 further comprising the steps of:

ordering said loaded one or more flows classified by CIX and DA or CIX only which are between said lowest index and said gMaxCixDaOnlyFix index for subsuming ordering; and

ordering said loaded one or more flows classified by DA only which are between said gMinDaOnlyFix index and said partition index for subsuming ordering.

12. The method of claim 9 wherein said first partition is divided into a first portion and a second portion, said first portion includes indices having highest priority ranging from said lowest index to a gMaxCixDaFix index and said second portion includes indices having lowest priority ranging from a gMinDaOnlyFix index to said partition index, said gMaxCixDaFix index being adjacent to said gMinDaOnlyFix index and a free entry being positioned above said gMinDaOnlyFix index in said second portion and further comprising the steps of:

(a) moving a flow classified by CIX and DA or CIX only at said gMaxCixDaFix to said free entry;

(b) assigning a flow classified by DA only to said gMaxCixDaFix index after step (a);

(c) setting the gMinDaOnlyFix index at the index assigned to said flow classified by DA only in step (b); and

(d) setting the gMaxCixDaFix entry immediately before the index assigned to said flow classified by DA only in step (b).

13. The method of claim 12 further comprising the steps of:

5 ordering said loaded one or more flows classified by DA only which are between said partition index and said gMinDaOnlyFix index for subsuming ordering; and
 ordering said loaded one or more flows classified by CIX and DA or CIX only between said gMaxCixDaFix index and said lowest index for subsuming ordering.

10 14. The method of claim 7 further comprising the steps of:

 maintaining a flow index space having entries corresponding to said TCAM; and
 assigning said flows classified by CIX and DA and CIX only to a CIX prefix tree.

15 15. The method of claim 7 further comprising the steps of:

 maintaining a flow index space having entries corresponding to said TCAM; and
 assigning said flows classified by DA only to a DA prefix tree.

16. The method of claim 7 further comprising the step of:

20 removing a flow in said TCAM by freeing up a corresponding said entry in said flow index space and invalidating a corresponding said TCAM entry.

17. The method of claim 1 wherein said predetermined order of said first partition has one of said TCAM entries with a longest prefix located at an index having highest priority followed by decreasing prefix values with a shortest prefix at an index
 25 having lowest priority.

18. The method of claim 3 wherein said second partition is divided into a first portion and a second portion, said first portion includes indices having highest priority ranging from a lowest index in said second partition to a GACL-partition index and said
 30 second portion includes indices having lowest priority ranging from said GACL-partition

index to said highest index, said ACL flows are assigned to said first portion and global access list flows are assigned to said second portion.

19. The method of claim 18 wherein said ACL flow specifies a range of source or destination ports, and further comprising the step of:

mapping said ACL flow to multiple flows in said TCAM wherein said mask for each of said multiple flows covers a portion of said range of source or destination ports for optimally determining a number of flows to cover said portion of said range.

20. The method of claim 18 further comprising the step of:

maintaining a flow index space having entries corresponding to said TCAM.

21. The method of claim 20 further comprising the steps of:

removing a flow in said TCAM by freeing up a corresponding said entry in said flow index space and invalidating a corresponding said TCAM entry; and

compacting said flows in said TCAM by moving each remaining said TCAM entry up by one index.

22. The method of claim 19 further comprising the steps of:

maintaining a flow index space having entries corresponding to said TCAM.

23. The method of claim 22 further comprising the step of:

removing said multiple flows in said TCAM by freeing up corresponding said entries in said flow index space and invalidating said corresponding entries in said

TCAM entry; and

compacting said flows in said TCAM by moving each remaining said TCAM entry up by one or more indices.

24. A method for classifying a plurality of data flows in a router comprising the steps of:

partitioning a ternary content addressable memory (TCAM) into at least a first partition and a second partition, said first partition includes indices having highest priority ranging from a lowest index to a partition index and said second partition includes indices having lowest priority ranging from a highest index to said partition index;

loading one or more first flow TCAM entries of a first of said plurality of data flows into said first partition in a predetermined order;

loading one or more second TCAM entries of a second of said plurality of data flows into said second partition in a predetermined order;

setting bit values of a corresponding mask for each of said first TCAM entries and said second TCAM entries such that bits of said respective first TCAM entries and said second TCAM entries are individually masked by said masks; and

comparing a prefix comprising predetermined packet header information of an incoming packet to said loaded one or more first TCAM entries and one or more second TCAM entries such that a matching said one or more first TCAM entries subsumes any matching said one or more second TCAM entries,

wherein said first plurality of data flows are MPLS or IP-VPN flows and said second plurality of data flows are policy based routing flows.

25. A method for classifying a plurality of data flows in a router comprising the steps of:

partitioning a ternary content addressable memory (TCAM) into at least a first partition and a second partition, said first partition includes indices having highest priority ranging from a lowest index to a partition index and said second partition includes indices having lowest priority ranging from a highest index to said partition index;

loading one or more first flow TCAM entries of a first of said plurality of data flows into said first partition in a predetermined order;

loading one or more second TCAM entries of a second of said plurality of data flows into said second partition in a predetermined order;

setting bit values of a corresponding mask for each of said first TCAM entries and said second TCAM entries such that bits of said respective first TCAM entries and said second TCAM entries are individually masked by said masks;

comparing a prefix comprising predetermined packet header information of an incoming packet to said loaded one or more first TCAM entries and one or more second TCAM entries such that a matching said one or more first TCAM entries subsumes any matching said one or more second TCAM entries;

maintaining a flow index space having entries corresponding to said TCAM; and

determining said predetermined order of said first TCAM entries and said predetermined order of said second TCAM entries in said flow index space before said steps of loading said one or more first TCAM entries.

26. A system for classifying a plurality of data flows in a router comprising:

means for partitioning a ternary content addressable memory (TCAM) into at least a first partition and a second partition, said first partition includes indices having highest priority ranging from a lowest index to a partition index and said second partition includes indices having lowest priority ranging from a highest index to said partition index;

means for loading one or more first flow TCAM entries of a first of said plurality of data flows into said first partition in a predetermined order;

means for loading one or more second TCAM entries of a second of said plurality of data flows into said second partition in a predetermined order;

means for setting bit values of a corresponding mask for each of said first TCAM entries and said second TCAM entries such that bits of said respective first TCAM entries and said second TCAM entries are individually masked by said masks; and

means for comparing a prefix comprising packet header information of an incoming packet to predetermined said loaded one or more first TCAM entries and one or more second TCAM entries,

wherein a matching said one or more first TCAM entries subsumes an matching said one or more second TCAM entries.

27. The system of claim 26 where said first plurality of data flows are MPLS or IP-VPN flows.

28. The system of claim 26 wherein said second plurality of data flows are policy based routing flows.

29. The system of claim 28 wherein said policy based routing flows are access control list (ACL) flows.

30. The system of claim 28 wherein said policy based routing flows are traffic manager flows.

31. The system of claim 26 further comprising:
means for maintaining a flow index space having entries corresponding to said TCAM; and
means for determining said predetermined order of said first TCAM entries and said predetermined order of said second TCAM entries in said flow index space.

32. The system of claim 26 wherein said MPLS or IP-VPN flows are classified by connection index (CIX) and destination address (DA), CIX only or DA only.

33. The system of claim 32 wherein said first partition is divided into a first portion and a second portion, said first portion includes indices having highest priority ranging from said lowest index to a gMaxCixDaFix index and said second portion includes indices having lowest priority ranging from a gMinDaOnlyFix index to said partition index and said flows classified by CIX and DA and CIX only are assigned to said first portion and said flows classified by DA only are assigned to said second portion.

34. The system of claim 33 wherein said flows classified by CIX and DA and said flows classified by CIX only are loaded into said first portion starting at said

gMaxCixDaFix index and said flows classified by said DA only are loaded in said second portion starting at said gMinDaOnlyFix index.

35. The system of claim 32 further comprising:

5 means for maintaining a flow index space having entries corresponding to said TCAM; and assigning said flows classified by CIX and CIX only to a CIX prefix tree.

36. The system of claim 32 wherein said first partition is divided into a first portion and a second portion, said first portion includes indices having highest priority
10 ranging from said lowest index to a gMaxCixDaFix index and said second portion includes indices having lowest priority ranging from a gMinDaOnlyFix index to said partition index, said gMaxCixDaFix index being adjacent to said gMinDaOnlyFix index and a free entry being positioned below said gMinDaOnlyFix index in said second portion and further comprising:

15 means for moving a flow classified by DA only at said gMinDaOnlyFix index to said free entry;

means for assigning a flow classified by CIX and DA or CIX only to said gMinDaOnlyFix index;

20 means for setting the gMaxCixDaFix index at the index assigned to said flow classified by CIX and DA or CIX only; and

means for setting the gMinDaOnlyFix index immediately after the index assigned to said flow classified by CIX and DA or CIX only.

37. The system of claim 36 further comprising:

25 means for ordering said loaded one or more flows classified by CIX and DA or CIX only which are between said lowest index to said gMaxCixDaOnlyFix index for subsuming ordering; and

means for ordering said loaded one or more flows classified by DA only which are between said gMinDaOnlyFix index and said partition index for subsuming ordering.

38. The system of claim 32 wherein said first partition is divided into a first portion and a second portion, said first portion includes indices having highest priority ranging from said lowest index to a gMaxCixDaFix index and said second portion includes indices having lowest priority ranging from a gMinDaOnlyFix index to said partition index, said gMaxCixDaFix index being adjacent to said gMinDaOnlyFix index and a free entry being positioned above said gMinDaOnlyFix index in said second portion and further comprising:

means for moving a flow classified by CIX and DA or CIX only at said gMaxCixDaFix to said free entry;

means for assigning a flow classified by DA only to said gMaxCixDaFix index;

means for setting the gMinDaOnlyFix index at the index assigned to said flow classified by DA only; and

means for setting the gMaxCixDaFix entry immediately before the index assigned to said flow classified by DA only.

39. The system of claim 38 wherein said first partition is divided into a first portion and a second portion, said first portion includes indices having highest priority ranging from said lowest index to a gMaxCixDaFix index and said second portion includes indices having lowest priority ranging from a gMinDaOnlyFix index to said partition index, said gMaxCixDaFix index being adjacent to said gMinDaOnlyFix index and a free entry being positioned above said gMinDaOnlyFix index in said second portion and further comprising:

means for moving a flow classified by CIX and DA or CIX only at said gMaxCixDaFix to said free entry;

means for assigning a flow classified by DA only to said gMaxCixDaFix index;

means for setting the gMinDaOnlyFix index at the index assigned to said flow classified by DA only; and

means for setting the gMaxCixDaFix entry immediately before the index assigned to said flow classified by DA only.

40. The system of claim 32 further comprising:

means for maintaining a flow index space having entries corresponding to said TCAM; and

means for assigning said flows classified by DA only to a DA prefix tree.

5 41. The system of claim 40 further comprising:

means for removing a flow in said TCAM by freeing up a corresponding said entry in said flow index space and invalidating a corresponding said TCAM entry.

10 42. The system of claim 26 wherein said predetermined order of said first partition has one of said TCAM entries with a longest prefix located at an index having highest priority followed by decreasing prefix values with a shortest prefix at an index having lowest priority.

15 43. The system of claim 28 wherein said second partition is divided into a first portion and a second portion, said first portion includes indices having highest priority ranging from a lowest index in said second portion to a GACL-partition index and said second portion includes indices having lowest priority ranging from said GACL-partition index to said highest index, said ACL flows are assigned to said first portion and global access list flows are assigned to said second portion.

20 44. The system of claim 43 wherein said ACL flow specifies a range of source or destination ports, and further comprising:

25 means for mapping said ACL flow to multiple flows in said TCAM wherein said mask for each of said multiple flows covers a portion of said range of source or destination ports.

45. The system of claim 43 further comprising:

means for maintaining a flow index space having entries corresponding to said TCAM.

30

46. The system of claim 45 further comprising:

means for removing a flow in said TCAM by freeing up a corresponding said entry in said flow index space and invalidating a corresponding said TCAM entry; and
 means for compacting said flows in said TCAM by moving each remaining said TCAM entry up by one or more indices.

5

47. The system of claim 44 further comprising:

means for maintaining a flow index space having entries corresponding to said TCAM.

10

48. The system of claim 47 further comprising:

means for removing said multiple flows in said TCAM by freeing up corresponding said entries in said flow index space and invalidating said corresponding entries in said TCAM entry; and

means for compacting said flows in said TCAM by moving each remaining said TCAM entry up by one or more indices.

15

49. A system for classifying a plurality of data flows in a router comprising:

means for partitioning a ternary content addressable memory (TCAM) into at least a first partition and a second partition, said first partition includes indices having highest priority ranging from a lowest index to a partition index and said second partition includes indices having lowest priority ranging from a highest index to said partition index; loading one or more first flow TCAM entries of a first of said plurality of data flows into said first partition in a predetermined order;

20

means for loading one or more second TCAM entries of a second of said plurality of data flows into said second partition in a predetermined order;

25

means for setting bit values of a corresponding mask for each of said first TCAM entries and said second TCAM entries such that bits of said respective first TCAM entries and said second TCAM entries are individually masked by said masks; and

means for comparing a prefix comprising predetermined packet header

information of an incoming packet to said loaded one or more first TCAM entries and

30

one or more second TCAM entries such that a matching said one or more first TCAM entries subsumes any matching said one or more second TCAM entries,

wherein said first plurality of data flows are MPLS or IP-VPN flows and said second plurality of data flows are policy based routing flows.

5

50. A system for classifying a plurality of data flows in a router comprising:

means for partitioning a ternary content addressable memory (TCAM) into at least a first partition and a second partition, said first partition includes indices having highest priority ranging from a lowest index to a partition index and said second partition includes indices having lowest priority ranging from a highest index to said partition index;

10

means for loading one or more first flow TCAM entries of a first of said plurality of data flows into said first partition in a predetermined order;

means for loading one or more second TCAM entries of a second of said plurality of data flows into said second partition in a predetermined order;

15

means for setting bit values of a corresponding mask for each of said first TCAM entries and said second TCAM entries such that bits of said respective first TCAM entries and said second TCAM entries are individually masked by said masks;

means for comparing a prefix comprising predetermined packet header information of in incoming packet to said loaded one or more first TCAM entries and one or more second TCAM entries such that a matching said one or more first TCAM entries subsumes an matching said one or more second TCAM entries;

20

means for maintaining a flow index space having entries corresponding to said TCAM; and

25

means for determining said predetermined order of said first TCAM entries and said predetermined order of said second TCAM entries in said flow index space before said steps of loading said one or more first TCAM entries.

51. An apparatus for classifying a plurality of data flows in a routing system comprising:

30

a ternary content addressable memory (TCAM);

a partitioning algorithm for partitioning said TCAM into at least a first partition and a second partition, said first partition includes indices having highest priority ranging from a lowest index to a partition index and said second partition includes indices having lowest priority ranging from a highest index to said partition index;

5 a loading algorithm for selecting a respective mask value to structure one or more first flow TCAM entries of a first of said data flows and one or more second flow TCAM entries and said respective mask values into said second partition; and

a search algorithm for performing an associative comparison of a prefix comprising predetermined packet header information of an incoming packet to said
10 loaded one or more first flow TCAM entries and one or more second TCAM entries of a first of said plurality of data flows into said first partition in a predetermined order such that a matching said one or more first TCAM entries subsumes an matching said one or more second TCAM entries;

15 52. The apparatus of claim 51 wherein said first plurality of data flows are MPLS or IP-VPN flows.

53. The apparatus of claim 51 wherein said second plurality of data flows are policy based routing flows.

20 54. The apparatus of claim 51 wherein said partitioning algorithm and said loading algorithm are implemented in flow index space having entries corresponding to said TCAM and said predetermined order of said first TCAM entries and said predetermined order of said second TCAM entries is represented in said flow index
25 space.

55. The apparatus of claim 52 wherein said MPLS or IP-VPN flows are classified by connection index (CIX) and destination address (DA), CIX only or DA only.

30 56. The apparatus of claim 55 wherein said first partition is divided into a first portion and a second portion, said first portion includes indices having highest priority

ranging from said lowest index to a gMaxCixDaFix index and said second portion includes indices having lowest priority ranging from a gMinDaOnlyFix index to said partition index and said flows classified by CIX and DA and CIX only are assigned to said first portion and said flows classified by DA only are assigned to said second
 5 portion.

57. The apparatus of claim 56 wherein said flows classified by CIX and DA and flows classified by CIX only are loaded into said first portion starting at a beginning of said TCAM and ending at said gMaxCixDaFix index and said flows classified by said
 10 DA only are loaded in said second portion starting at said gMinDaOnlyFix index.

58. The apparatus of claim 52 further comprising:
 means for maintaining a flow index space having entries corresponding to said
 TCAM; and
 15 means for assigning said flows classified by CIX to a CIX prefix tree.

59. The apparatus of claim 52 further comprising:
 means for maintaining a flow index space having entries corresponding to said
 TCAM; and
 20 means for assigning said flows classified by DA only to a DA prefix tree.

60. The apparatus of claim 59 further comprising:
 means for removing a flow in said TCAM by freeing up a corresponding said
 entry in said flow index space and invalidating a corresponding said TCAM entry.
 25

61. The apparatus of claim 52 wherein said predetermined order of said first partition has one of said TCAM entries with a longest prefix located at an index having highest priority followed by decreasing prefix values with a shortest prefix at an index having lowest priority.
 30

62. The apparatus of claim 53 wherein said second partition is divided into a first portion and a second portion, said first portion includes indices having highest priority ranging from a lowest index in said second partition to a GACL-partition index and said second portion includes indices having lowest priority ranging from said GACL-partition index to said highest index in said second partition, said ACL flows are assigned to said first portion and global access list flows are assigned to said second portion.

63. The apparatus of claim 62 wherein said ACL flow specifies a range of source or destination ports, and further comprising:

10 means for mapping said ACL flow to multiple flows in said TCAM wherein said mask for each of said multiple flows covers a portion of said range of source or destination ports.

64. The apparatus of claim 62 further comprising:

15 means for maintaining a flow index space having entries corresponding to said TCAM.

65. The apparatus of claim 64 further comprising:

20 means for removing a flow in said TCAM by freeing up a corresponding said entry in said flow index space and invalidating a corresponding said TCAM entry; and means for compacting said flows in said TCAM by moving each remaining said TCAM entry up by one index.

66. The apparatus of claim 62 further comprising:

25 means for maintaining a flow index space having entries corresponding to said TCAM.

67. The apparatus of claim 62 further comprising:

30 means for removing said multiple flows in said TCAM by freeing up corresponding said entries in said flow index space and invalidating said corresponding entries in said TCAM entry; and

means for compacting said flows in said TCAM by moving each remaining said TCAM entry up by one or more indices.